

JAN 11 2001

REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-01-

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Avenue, Washington, DC 20540-6001, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

Maintaining
tions for
Office of

0042

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 01/07/01		3. REPORT TYPE Final (November 1997 -- November 2000)	
4. TITLE AND SUBTITLE Research Issues in the Next Generation Group Communication Services.				5. FUNDING NUMBERS F49620-98-1-0070	
6. AUTHOR(S) Shivakart Mishra					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Wyoming 9th & Iverson Street Laramie, WY 82071-3355.				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/PKC 801 North Randolph Street, Room 732 Arlington, VA 22203-1977.				10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unlimited. DISTRIBUTION STATEMENT A Approved for Public Release Distribution Unlimited					
13. ABSTRACT (Maximum 200 Words) The main objective of the proposed research was to investigate five important practical issues in the design of high performance, heterogeneous group communication services for constructing the next generation defense applications. These issues are: (1) Fail Awareness : providing fail awareness property in group communication services under a <i>timed asynchronous distributed system model</i> , (2) Heterogeneity : handling heterogeneity in computing components and communication networks in the design and implementation of group communication services, (3) High Speed Communication : exploiting high-speed communication networks such as ATM, Myrinet, and Internet2 in the design and implementation of group communication services, (4) Versatility : improving the versatility of group communication services, and (5) High Performance : providing high performance under several different operating conditions.					
14. SUBJECT TERMS Group Communication, CORBA, Agent, Communication				15. NUMBER OF PAGES 11	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified		20. LIMITATION OF ABSTRACT UL	

OBJECTIVES

The main objective of the proposed research was to investigate the following five important practical issues in the design of high performance, heterogeneous group communication services for constructing the next generation defense applications.

1. **Fail Awareness:** Providing fail awareness property in group communication services under a *timed asynchronous distributed system model*.
2. **Heterogeneity:** Handling heterogeneity in computing components and communication networks in the design and implementation of group communication services.
3. **High Speed Communication:** Exploiting high-speed communication networks such as ATM, Myrinet, and Internet2 in the design and implementation of group communication services.
4. **Versatility:** Improving the versatility of group communication services.
5. **High Performance:** Providing high performance under several different operating conditions.

STATUS OF EFFORT

All five issues have been investigated as a part of this research and some significant results were obtained. This was done through the continuation of one project, the timewheel group communication service, funded by an earlier AFOSR grant and the initiation of two new projects, CORBA group communication service and DaAgent.

CORBA Group Communication Service

This project investigated the design, implementation, and performance of a group communication service implemented using the Common Object Request Broker Architecture (CORBA). In the process, this project addressed two important issues in the design and implementation of group communication services: heterogeneity and high performance. The key questions addressed are: (1) how practical is it to implement a group communication service using CORBA?, (2) what is the performance overhead of such an implementation?, (3) what are the sources of this performance overhead?, and (4) are the current CORBA specifications sufficient for implementing group communication services using the service approach? The timewheel atomic broadcast protocol and a three-round, majority agreement group membership protocol were implemented in three ways: (1) using UDP socket interface, (2) using CORBA, and (3) using CORBA and UDP socket interface. The main conclusion of this project is that the current CORBA technology can be used to provide heterogeneity properties to group communication services at some performance cost. A naive implementation of a group communication service in CORBA incurs substantial performance overhead. So, it is very important to tune various group communication protocol implementations in CORBA to reduce this performance overhead. In particular, the lengths of various temporary buffers and the timers used by the group communication service must be tuned very carefully.

DaAgent

The DaAgent system provides an agent-based computing paradigm for dependable computing over the Internet. This system is written in Java, allows creation and migration of Java agents, and runs in a heterogeneous, distributed computing environment such as the Internet. This project addressed three important issues in the design and implementation of dependable distributed system: fail awareness, heterogeneity, and performance. There are four unique features of DaAgent. First, it has been designed specifically for experimenting with dependability protocols (security and fault-tolerance) in an agent-based computing system. The main focus in the design of DaAgent has been on modularizing various implementation modules to

facilitate experimentation. It is possible to redesign and reimplement a single module within DaAgent without affecting any other module. Second, DaAgent is the first agent-based computing system that addresses fault-tolerance issues such as node and communication failures. Third, the DaAgent system is the first system that has focussed at the two dependability issues, namely security and fault tolerance, together. Finally, the DaAgent system provides extensive support for interagent communication and synchronization. In particular, it provides eight different models of interagent communication and synchronization.

Timewheel Group Communication Service

A new group communication service called the *timewheel group communication service* has been designed and implemented. An extensive performance measurement has shown that this service provides one of the best overall performance under several different operating conditions. This service provides four unique characteristics:

1. this is the first group communication service designed specifically for a *timed* asynchronous distributed system model. This system allows the construction of dependable protocols that specify what outputs and state transitions should occur in response to inputs as well as a real-time interval in which these outputs and state transition will occur.
2. the three protocols that comprise the timewheel group communication service, namely clock synchronization, atomic broadcast, and group membership, are fail aware. This means that they provide mechanisms by which a user is informed of the failures when these protocols may not be able to provide real-time guarantees
3. the timewheel group communication service supports multiple semantics simultaneously: three order semantics—*unordered*, *total ordered* and *time ordered*, three atomicity semantics—*weak atomicity*, *strong atomicity* and *strict atomicity*, and one termination semantic. These result in a total of nine group communication semantics. A user can dynamically choose one of these nine semantics while broadcasting an update.
4. the strict atomicity together with the time order semantics provided by the timewheel group communication system preserves causal order between updates that may arise due to “hidden” channels. This is a first attempt in preserving causal order arising due to “hidden” channels in an asynchronous system.

ACCOMPLISHMENTS/NEW FINDINGS

The research addressed five important practical issues in the design and implementation of group communication services. This has resulted in several significant findings:

1. CORBA can be used to build high performance group communication services that operate in a heterogeneous distributed computing environment. A naive implementation incurs substantial performance overhead due to CORBA. However, this performance overhead can be reduced to some extent by choosing appropriate implementation techniques, and carefully tuning various group communication parameters such as buffer sizes and timer lengths.
2. Security and fault-tolerance protocols interact with one another extensively in an agent-based computing system. As a result, these protocols must be designed together to ensure a high degree of dependability.
3. A new synchronization mechanism called location synchronization has been proposed. This synchronization mechanism allows two or more agents to meet at a specific location in the Internet.
4. It is possible to provide multiple group communication semantics ranging from very weak to very strong in a group communication service and still preserve good overall performance of such a service.
5. Real time guarantees can be provided in a group communication service even when it is implemented in a non-synchronous distributed computing environment.

PERSONNEL SUPPORTED

Faculty: Shivakant Mishra

Post-Docs: None

Graduate Students: Bozheng Yang, Xuyang Jiang, Guming Xing, Lan Fei, Yanjun Huang, Harshavardhan Kuntur, Xiao Lin, and Peng Xie.

Other (Undergraduate Students): William Kirkwood and Nija Shi.

PUBLICATIONS

Journal Publications

On Group Communication Support in CORBA (with L. Fei, X. Lin and G. Xing). *IEEE Transactions on Parallel and Distributed Systems*. To appear in January 2001.

Newsmonger: A Technique to Improve the Performance of Atomic Broadcast Protocols (with S. Kuntur). *The Journal of Systems and Software*. Vol. 55, No. 2. (January 2001).

Low Level Support for Implementing Group Communication Services (with F. Cristian). *ISCA International Journal of Computers and Their Applications*. Vol. 5, No. 4 (Dec 1998).

An Evaluation of Flow Control in Group Communication (with L. Wu). *IEEE/ACM Transactions on Networking*. Vol. 6, No. 5 (Oct 1998).

High Performance Asynchronous Atomic Broadcast (with F. Cristian and G. Alvarez). *Distributed Systems Engineering Journal*. Vol. 4, No. 2 (June 1997), 109–128.

Units of Computation in Fault-Tolerant Distributed Systems (with M. Ahuja). *Journal of Parallel and Distributed Computing*, Vol. 40, No. 2 (February 1997), 194–209.

Refereed Conference Publications

Interagent Communication and Synchronization in DaAgent (with P. Xie). *Proceedings of the 21st IEEE International Conference on Distributed Computing Systems*, Phoenix, AZ (April 2001). To appear.

Statistical Analysis of Security Violations in Computer Software (with S. Gui). *Proceedings of the SCA 16th International Conference on Computers and Their Applications*, Seattle, WA (March 2001). To appear.

Design, Implementation and Performance Evaluation of a High Performance CORBA Group Membership Protocol (with X. Lin). *Proceedings of the 7th IEEE International Conference on High Performance Computing*, Bangalore, India (December 2000).

Programming Environment of DaAgent (with Y. Huang, and H. Kuntur). *Proceedings of the 12th IASTED International Conference Parallel and Distributed Computing and Systems*, Las Vegas, NV (November 2000).

Fault Tolerance in Agent-Based Computing Systems (with Y. Huang). *Proceedings of the 13th ISCA International Conference on Parallel and Distributed Computing Systems*, Las Vegas, NV (August 2000).

Models for Interagent Communication and Synchronization (with P. Xie). *Proceedings of the 2000 International Conference on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, NV (June 2000).

Specifying and Checking Fault-tolerant Agent-based Protocols using Maude (with J. van Baalen and J. Caldwell). *Proceedings of the 1st Goddard Workshop on Formal Approaches to Agent-Based Systems*, Greenbelt, MD (April 2000).

Replicated Revision Control System (with J. Ward). *Proceedings of the 11th IASTED International Conference on Parallel and Distributed Computing and Systems*, Cambridge, MA (November 1999).

Performance Evaluation of a CORBA Group Communication Service (with L. Fei and G. Xing). *Proceedings of the 29th IEEE International Symposium on Fault-tolerant Computing*, Madison, WI (June 1999).

Transparent TCP/IP based Replication (with C. Fetzer). Fastabstract. *Proceedings of the 29th IEEE International Symposium on Fault-tolerant Computing*, Madison, WI (June 1999).

DaAgent: A Dependable Mobile Agent System (with Y. Huang, and H. Kuntur). Fastabstract. *Proceedings of the 29th IEEE International Symposium on Fault-tolerant Computing*, Madison, WI (June 1999).

Providing Fault Tolerance to Mobile Intelligent Agents (with X. Jiang and B. Yang). *Proceedings of the ISCA 8th International Conference on Intelligent Systems*, Denver, CO (June 1999).

Design and Implementation of an Availability Management Service (with G. Pang). *Proceedings of the ICDCS Workshop on Middleware*, Austin, TX (June 1999).

Improving Performance of Atomic Broadcast Protocols Using the Newsmonger Technique (with S. M. Kuntur). *Proceedings of the Seventh IFIP International Working Conference on Dependable Computing for Critical Applications*, San Jose, CA (January 1999).

Constructing Applications Using the Timewheel Group Communication Service. *Proceedings of the 1998 International Conference on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, NV (July 1998).

The Timewheel Group Membership Protocol (with C. Fetzer and F. Cristian). *Proceedings of the 3rd IEEE Workshop on Fault-tolerant Parallel and Distributed Systems*, Orlando, FL (April 1998).

Thread-based vs Event-based Implementation of a Group Communication Service (with R. Yang). *Proceedings of the 12th IEEE International Parallel Processing Symposium & 9th IEEE Symposium on Parallel and Distributed Processing*, Orlando, FL (April 1998).

INTERACTIONS/TRANSITIONS

Participation/Presentations At Meetings, Conferences, Seminars, Etc.

1. The 7th IEEE International Conference on High Performance Computing, Bangalore, India (December 2000).
2. The 12th IASTED International Conference Parallel and Distributed Computing and Systems, Las Vegas, NV (November 2000).
3. The 13th ISCA International Conference on Parallel and Distributed Computing Systems, Las Vegas, NV (August 2000), (Session Chair).
4. The 2000 International Conference on Parallel and Distributed Processing Techniques and Applications, Las Vegas, NV (June 2000).
5. The 1st Goddard Workshop on Formal Approaches to Agent-Based Systems, Greenbelt, MD (April 2000).
6. The 11th IASTED International Conference on Parallel and Distributed Computing and Systems, Cambridge, MA (November 1999), (Session Chair).
7. The 29th IEEE International Symposium on Fault-tolerant Computing, Madison, WI (June 1999).
8. The IEEE International Conference on Distributed Computing Systems, Austin, TX (May 1999).
9. The Seventh IFIP International Working Conference on Dependable Computing for Critical Applications, San Jose, CA (January 1999).
10. The 1998 International Conference on Parallel and Distributed Processing Techniques and Applications, Las Vegas, NV (July 1998), (Session Chair).
11. The 3rd IEEE Workshop on Fault-tolerant Parallel and Distributed Systems, Orlando, FL (April 1998).

Consultative And Advisory Functions To Other Laboratories And Agencies None.

Transitions None.

NEW DISCOVERIES, INVENTIONS, OR PATENT DISCLOSURES

1. CORBA can be used to build high performance group communication services.
2. To ensure overall dependability of agent-based computing systems, security and fault-tolerance protocols must be designed together.
3. Location synchronization is a very useful synchronization technique in mobile computing applications.
4. It is possible to provide some real-time guarantees in an asynchronous distributed computing environment.

HONORS/AWARDS

Researcher of Extraordinary Merit Award. University of Wyoming, 1999.